Founsure 1.0: An Erasure Code Library with Efficient Repair and Update Features

Şuayb Ş. Arslan

MEF University, Maslak, Istanbul, Turkey

arXiv:1702.07409v1 [cs.IT] 23 Feb 2017

Abstract

Founsure is an open-source software library, distributed under LGPLv3 license and implements a multi-dimensional graph-based erasure coding entirely based on fast exclusive OR (XOR) logic. Its implementation uses compiler optimizations to generate the right assembly for the given SIMD-enabled architectures. Founsure 1.0 supports a variety of features that shall find interesting applications in modern data storage as well as communication and computer network systems which are becoming hungry in terms of network bandwidth, computational resources and average consumed power. In particular, Founsure erasure code provides a three dimensional design space i.e., computation complexity, coding overhead and repair bandwidth to meet the requirements of modern distributed data storage and processing systems in which the data needs to be protected against device/hardware failures.

Keywords: Distributed Storage, Erasure Coding,, Fountain Coding, Single Instruction Multiple Data (SIMD), Erasure coding, openMP, Reliability.

1. Motivation and significance

Erasure coding is used for fault tolerance and providing required reliability and availability of data in distributed data storage systems [1]. As the modern data storage systems evolved to have different requirements, the set of constraints on the design of erasure codes has changed. For instance, previous research on erasure codes mostly focused on optimizing the coding overhead i.e., minimization of storage space consumption for the given target data reliability [2], [3]. Moreover, some of these designs considered pure XOR operation to both provide required reliability and efficient computation [4]. More recently, locally repairable codes have attracted attention due to their efficient utilization of network resources and eventually achieve better data reliability [5]. Besides the proprietary implementations of erasure codes, many open source implementations based on different mathematical operations are available online [6, 7]. Main objective of these studies is to provide overhead optimal and fast erasure coding libraries without any consideration of the available network and computational resources and repair efficiency in a distributed setting.

Founsure 1.0 utilizes three-dimensional bipartite graph to construct multi-functional erasure code. The design space includes computation complexity, coding overhead and repair bandwidth as different dimensions of optimization. If one prefers to have an overhead optimal design with good performance, then it might be sufficient to use the most recent Jerasure 2.0 [8] erasure coding library by Dr. Plank and Dr. Greenan, now fully supported by *RedHat Ceph* Community [9]. On the other hand, the main objective of Founsure 1.0 is to provide different points in three dimensional design space based on the requirements of the storage applications through a set of parameter configurations. Founsure 1.0 library is successfully adapted and used together with a deduplication engine more recently and resulted in a conference paper publication [10].

The encoding process of Founsure 1.0 begins with a two-dimensional simple bipartite graph which leads to non-systematic fountain-like code [12]. The degree distribution of this code is specially selected to meet a good trade-off point between computation complexity, coding overhead and repair bandwidth. Version 1.0 supports Luby's Robust Soliton Distribution (RSD) [13] as well as Shokrollahi's Finite max-degree distribution [14] by default. However, we note that these distributions are optimized for minimum coding overhead only. On top of this bipartite graph, encoding engine generates check nodes for "data-only" (refered as check #1), "data & coding" (refered as check #2) and "coding-only" (refered as check #3) chunks/symbols. These check nodes can be represented in three dimensional space and shall be used to provide advanced decoding, repair and update features of Founsure. Throughout the document, we use nodes, chunks and symbols interchangeably.

Founsure 1.0 uses Belief Propagation (BP) [11] (a.k.a. message passing) algorithm to resolve or decode the user data, to repair the encoded data or update the encoded data. In version 1.0, the user data does not appear at the output in pure/plain format. In otherwords, one cannot read off data from the encoder output without any further processing. Therefore, using Founsure 1.0 encoding, user data can be considered encrypted automatically. We use pseudo-random number generators and seeds to generate graph connections of the underlying bipartite graph. So without the key/s (referred as seed/s in our implementation), there is no way to recover original user data. Therefore, Founsure 1.0 software package also provides a user configurable light weight built-in encryption feature.

This paper will briefly describe the set of functionalities provided with Founsure 1.0 and the details of the software architecture. The source code, as well as a comprehensive user guide, few test results and related documentation is available from *http://www.suaybarslan/founsure.html*.

2. Software description and architecture

2.1. Software Functionalities

Founsure 1.0 has the following three executable main functions that achieve four important functionalities.

- **founsureEnc**: Encoder engine that generates s number of data chunks under a local *Coding* directory and a metadata file that include information about the file as well as the coding parameters.
- **founsureDec**: Decoder engine that requires a local *Coding* directory, a valid file name and an associated metadata to run multiple BP passes.
- **founsureRep**: Repair engine that
  - fixes/repairs one or more coding chunks if they are erased, corrupted or flagged as unavailable.
  - generates extra coding chunks if update is requested. Update is a useful functionality particularly if data reliability is decreased/degraded over time.

There are also utility functions of Founsure 1.0 which are provided to help users to make correct design choices on reliability, complexity and storage space efficiency. We also use utility functions to trigger update functionality as will be demonstrated later. Utility functions do not directly process user data. Version 1.0 currently supports two utility functions as listed below.

- **simDisk**: This function can be used to exhaust all the possible combinations of disk failures for a given set of coding parameters. In other words, this function checks whether the provided coding parameters are sufficient to achieve a user defined reliability goal. Therefore, running this function can help us design target-policy erasure codes.

- **genChecks**: This utility function is crucial for two different functionalities: (1) fast/efficient repair and (2) smooth update. In the repair process, it generates two types of checks: Check #2 and Check #3 and writes them to a <testfile>_check.data file in a format described within this document. In case of update, it modifies the metadata file as well as <testfile>_check.data file so that the coding chunks can be updated by running **founsureRep** function.

Next, we provide the details of Founsure 1.0 encoding and decoding operations, particularly the implementation details of **FounsureEnc** and **FounsureDec** functions. For theoretical details, we refer the reader to the reference document [15].

2.2. Implementation details of Encoding/Decoding

In graph terminology, nodes (sometimes referred to as equations) are represented as graph vertices and node relationships are represented by edges of the graph. There are three types of nodes in a 3-D bipartite graph; data nodes, coding nodes and check nodes. The coding nodes represent a set of linear combinations of data nodes generated through a predetermined mathematical function such as Exclusive Or (XOR) logic operation. The final node type is known as check nodes. These nodes represent all the local sets of data and coding nodes for which a certain mathematical relationship is satisfied such as even/odd *parity*. A simple mathematical function that is used by Founsure 1.0 is XOR operation that utilizes multiple data blocks and generate a single block of information. We use f flag to indicate the file name, k to indicate the total number data nodes where b of these are the original user data nodes, n to indicate the total number of coding nodes and t to indicate the number of bytes to store per node.

In **FounsureEnc** function, data file with *filesize* bytes is partitioned into multiple $b \times t$ bytes and each partition is encoded independently as shown in Figure 1. As of version 1.0, *partition coupling* is not supported between distinct partitions. This technique is currently under investigation and might have interesting performance improvements to our design/implementation in analogy to spatially-coupled LDPC codes [16]. If *filesize* is not a multiple of $b \times t$ bytes, then we use zero padding to make *filesize* a multiple. **FounsureEnc** also checks whether $s|n$. If not, the least biggest n is selected automatically to satisfy $s|n$.

Encoding proceeds as follows. First, a memory space of worth $(k+n)t$ bytes is allocated. Next, check # 1 equations are generated by the efficient array LDPC encoding [18], [17]. An extra $k-b$ chunks are created to make up a total of k chunks of data. This process is shown as (1) in Figure 1. Later, a total of n coding chunks are generated from the whole set of k data chunks based on a fountain-like code with "FiniteDist" degree and pseudo random selection distributions. This process is shown as (2) in Figure 1. Finally, n coding chunks are distributed (striped) equally across distinct output files for allocation on s number of drives. We repeat this process for each data partition in a looped subprocess and append coding chunks at the end of the corresponding output files. For a given *<filename>.ext* file, we use *<filename>_ disk0..0i.ext* to refer to the ith output file. The number of zeros that appear in the name of output files is controlled by the "parameter.h" variable DISK_INDX_STRNG_LEN.

In our implementation, we have distinct object definitions for encoding, decoding, and repair operations. These objects have extension "*Obj" and include all the common parameters in their object fields. For instance, both encoding and/or decoding functions accept *EncoderObj* and/or *DecoderObj* constructs as inputs. Similarly, b and k variables can be accessed using standard way *EncoderObj.sizesb* and *EncoderObj.sizek*.

Each encoding/decoding object is associated with an initial seed value (*EncoderObj.seed*, default value is 1389488782) from which other seed values as well as the local sets of data chunks are pseudo-randomly created. Each coding chunk within *EncoderObj* and *DecoderObj* has their own unique ID. These IDs are used to identify the coding chunks which might be erased.

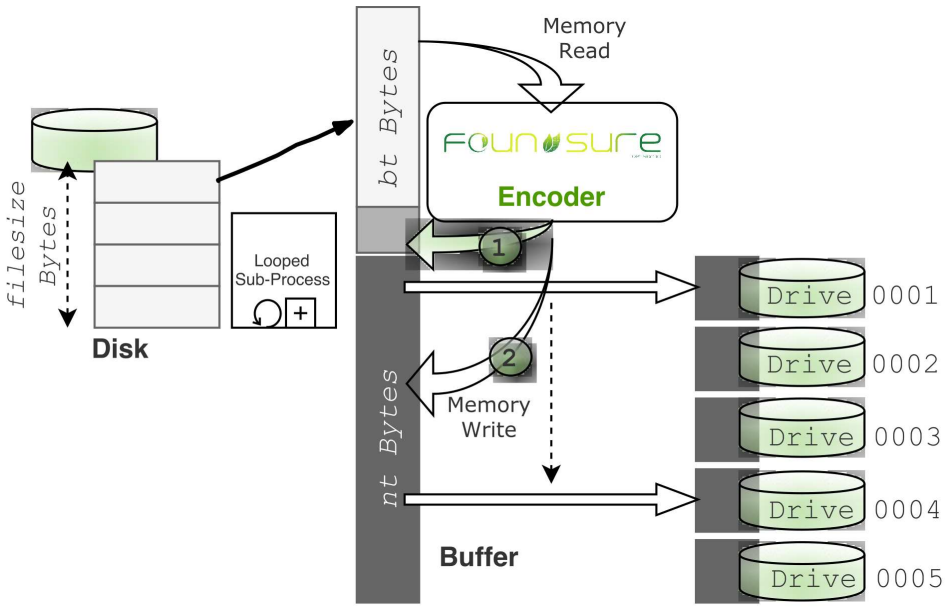

Figure 1: **FounsureEnc** allocates around $(n+k)t$ bytes of buffer and generates coding chunks and write it to shaded area on memory before the they are written to distinct drives. The file on disk is striped and processed in a looped subprocess as shown. The number of stripes is stored in *readin* variable and written to file metadata file.

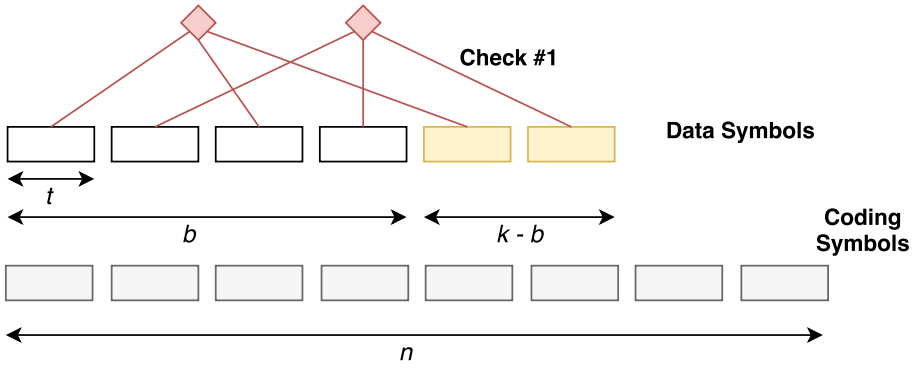

Figure 2: Founsure symbols are t bytes each and a precode is applied to find Check #1 equations. We add $k-b$ extra symbols to satisfy check equations as shown. Founsure encoding engine generates n coding chunks from k chunks as shown.

The seed value is used by the pseudorandom generator to produce a sequence of integers which form the basis of coding chunk degree number assignment as well as the selected data chunks for coding chunk computations. These numbers are stored as part of the object and can be regenerated using the same initial seed. Let us assume we have s number of output files, then we use $EncoderObj.seed + i$ as the seed of the $i-th$ output file with $0 \leq i < s$.

Coding symbols choose degrees first (from an appropriate degree distribution $\Omega(x)$) and that degree is used to select the number of data symbol neighbors to be involved in computation. The degree distribution $\Omega(x)$ is typically selected to minimize the coding overhead. For instance the following degree distribution is proposed for Raptor codes [14]

$$\Omega(x) = 0.007969x + 0.49357x^2 + 0.16622x^3 + 0.072646x^4 + 0.082558x^5 \\ +0.056058x^8 + 0.037229x^9 + 0.05559x^{19} + 0.025023x^{64} + 0003135x^{65}$$

However, Founsure does not necessarily minimize overhead. It may optimize overhead, repair bandwidth and complexity at the same time. We recommend to choose degree distributions

that will give us a good trade-off point between these three objectives. Although there is no optimal point for all applications, Founsure is designed to be highly configurable to fit in different requirements of modern storage systems.

We run **FounsureDec**, when we want to collect a subset of output data files and recover the original data file. Decoder is based on BP algorithm a summary of which is provided in **Algorithm 1**. BP function admits **DecoderObj** and the generator matrix of the code $\mathbf{B} \in \mathbb{F}_2^{k \times n}$. The decoder utilizes the information contained in metadata file to generate (prepare) the contents of *DecoderObj*. It works in a similar fashion to **FounsureEnc** i.e., it reads the striped coding chunks, runs BP algorithm at most twice (once for the outer graph code and if need be, additional one for the inner Array LDPC precode) and recovers the bt bytes at each turn. Finally, these bytes are written to decoded/recovered data file by calling standard kernel I/O commands.

Having all computation based on pseudorandomly selected chunks and carrying out these computations solely in terms of simple XOR logic has the cost of making the code non-optimal in terms of overhead (tough it might be near-optimal). If n coding symbols are distributed over s drives and when one of the drives fail, a subset of coding symbols are lost. In order to find what fraction of f-failure combinations can be tolerated, we provide a utility function **simDisk**.

2.3. Check Equations and Repair Process

We have three types of check nodes as will be discussed next.

Checks #1: These check equations are defined by the precode of the Founsure (for 1.0, we selected an Array LPDC code family [18]). Based on the selection of good precodes, their mathematical and coding parameter selections etc., the graph connections are automatically determined. Founsure 1.0 includes a precode support based on a binary array LDPC code and future releases will include external precode support which can be provided by the user using a preformatted input file. Please see precoding section to find more information about the construction of these check equations.

Checks #2: This type of check equations are generated by **Algorithm 3**. One of the special features of these checks is that only one neighbor is selected from the data nodes and the rest of the check terms are from the coding nodes. This special feature can be used to partially decode the data without running the complete decoder and unnecessarily reconstruct the unrequested parts of the data. An application of this could be multimedia in which the Region of Interest (RoI) can be directly reconstructed using this type of check equations.

Checks #3: These check equations are generated by **Algorithm 3**. These checks generate the local groups based on the coding nodes. These checks are primarily used to repair the erased, unavailable or unresponsive coding nodes in case of failures.

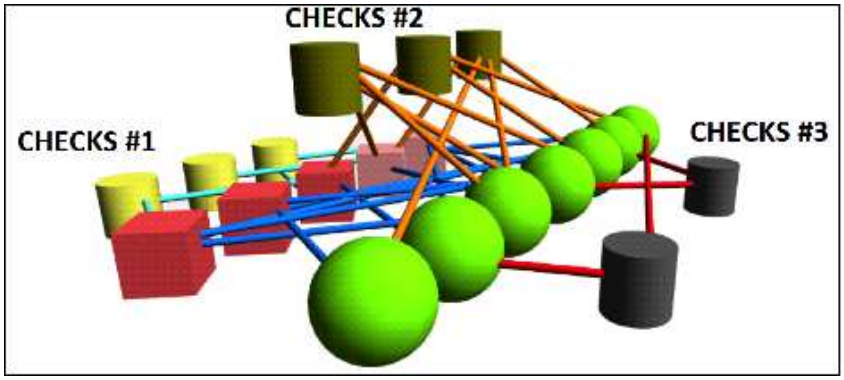

Figure 3: Founsure graph code is three dimensional bipartite graph with three sets of check equations.

Algorithm 1 Array LDPC Checks (Check # 1)

procedure CHECK1(b, k) ▷ Inputs: b, k
  $\mathbf{L}^{(1)} \leftarrow \mathbf{0}_{k \times k-b}$ ▷ Initial zero matrix
  $p \leftarrow$ largest_prime_factor(k)
  $kk \leftarrow k/p$
  $jj \leftarrow kk - b/p$
  for $j = 0; j < jj; j++$ **do**
    for $i = 0; i < p; i++$ **do**
      for $m = 1; m < kk - jj + 1; m++$ **do**
        $\mathbf{l}^{(1)}_{i+jp,m-1} \leftarrow kkp - (jj - j + m - 1)p - (m(jj - j - 1) - i - 1 + p) \pmod p - 1$
  return $\mathbf{L}^{(1)}$ ▷ $\mathbf{L}^{(1)}$: Check #1 sets.

2.4. Precoding - Generation of Check #1 equations

A (b, k, n) Founsure code takes b data chunks (bt bytes) and initially generates $k - b$ Check #1 parity chunks based on the binary array LDPC codes [18]. This special choice of array LDPC codes enable efficient encoding operation (linear with blocklength) and improves the complexity performance of the overall Founsure code.

The procedure outlined above uses a generic function largest_prime_factor(.) which chooses the largest prime factor of the argument. The rate of the array LDPC is defined as $r_{LDPC} = b/k$. The user can choose any k, n and r_{LDPC} and hence we can calculate $b = \lfloor kr_{LDPC} \rfloor$. Let $p =$ largest_prime_factor(k), we may not be able to find $\lfloor kr_{LDPC} \rfloor / p$ to be an integer. We can definitely use floor function to get an estimate of jj. However, the array LDPC code performance is heavily dependent on kk and jj values and there is no array LDPC code for all (k, r_{LDPC}) pairs. If (jj, kk) pair are small, the code performance is observed to be pretty bad. For this reason, we provide an algorithm that reasonably chooses a good performing LDPC code and satisfies (within some error margin) the user provided parameters k, n and r_{LDPC} at the same time.

Let us define the following system parameters before we formally give the algorithm that determines the closest good performing LDPC code for the user provided parameters k, n and r_{LDPC}. These system parameters are defined in "parameter.h" file.

- DIFF_TH: Allowed error threshold between the estimated and user provided b values.
- RRATE_TH: Allowed error threshold between the estimated and user provided precode rate.
- RED_BYTE_TH: Allowed redundant zero bytes to be appended at the end of the file for parameter consistency.
- RAND_WIN_MAX: Random number search window maximum value.
- RAND_WIN_MIN: Random number search window minimum value.
- ARRAY_MIN_JJ: Minimum Array LDPC "jj" parameter value.
- ARRAY_MIN_KK: Minimum Array LDPC "kk" parameter value.
- TRIES_TH: Threshold on the number of tries before incrementing DIFF_TH and RRATE_TH.
- DELTA_DIFF_TH: Step size increment for DIFF_TH
- DELTA_RRATE_TH: Step size increment for RRATE_TH.

Algorithm 2 Adjust Parameters w/ Precode (APP)

procedure APP(r_{LDPC},*filesize*,*b*,*t*)
  $\hat{b} \leftarrow b, b \leftarrow 0, k \leftarrow 1, kk \leftarrow 1e9, jj \leftarrow 1e9, p \leftarrow 2, blocks \leftarrow 1, tries \leftarrow 0, iter \leftarrow 0$
  while ($|b - \hat{b}|$ > DIFF_TH || |b/k - r_{LDPC}| > RRATE_TH || $kk > p$ || $jj > p$ || $kk < jj$ || |*blocks***t***b* - *filesize*| > RED_BYTE_TH **do**
    $k \leftarrow \lfloor \hat{b}/r_{LDPC} \rfloor + rand()$ mod (RAND_WIN_MAX + RAND_WIN_MIN) - RAND_WIN_MIN
    $p \leftarrow$ largest_prime_factor(k)
    $kk \leftarrow k/p, jj \leftarrow kk - \hat{b}/p$
    if $jj <$ ARRAY_MIN_JJ **then**
      $jj \leftarrow 2$
    $b \leftarrow (kk - jj)p$
    if $b > 0$ **then**
      $blocks \leftarrow filesize$ / tb + 1
      $tries++$
      if $tries >$ TRIES_TH **then**
        DIFF_TH ← DIFF_TH + DELTA_DIFF_TH
        RRATE_TH ← RRATE_TH + DELTA_RRATE_TH
        $tries \leftarrow 0$
    $iter++$
    if $iter > 1e7$ **then**
      print error and exit.
  $blocks \leftarrow filesize$ / (tb) + 1
  $redundantzeros \leftarrow bt \times blocks$ - *filesize*
  return $b, k, redundantzeros$

Next, we provide the algorithm that returns the estimated values of b, k and redundant number of zeros ($redundantzeros$). The algorithm admits four inputs, namely r_{LDPC}, $filesize$, b and t. Initial values of system parameters shall be set by "parameter.h" file and are changed locally within the function implementing **Algorithm 2**.

2.5. Information for Efficient Repair

Check #3 nodes are the most suitable node type for the repair process since it directly defines the relationship between the coded chunks. It is not hard to see having more of these type of nodes (created independently or dependently) gives alternative ways of repairing a given node in case of different combinations of node failures occur in the network. In other words, *the more of these type of checks, the more potential we have for the regeneration of the lost coded chunks*. With regard to this observation, we can use two techniques to increase the number of check #3 type information based on check #1 and check #2 equations:

- Identify the coded nodes with degree one (say we have M of those). Identify their data node neighbors. There are M check #2 type check equations that connect these data nodes with the coded nodes. Since the corresponding M coded nodes carry the same information, we can use these check #2 type check equations as check #3 check equation. Note that since check #2 and check #3 equations are derived from the same Founsure base graph, this technique is likely to generate already existent local groups or dependent local groups for coded nodes.

- Note that check #1 is user defined although subject to a predefined structure. This defines local groups over the data nodes. Since each data node is linked to local groups of coded nodes through check #2, we can use this relationship to derive check #3 type check

Algorithm 3 Check Generation Algorithm (CGA)

procedure CGA($\mathbf{B}$) ▷ Input matrix $\mathbf{B} \in \mathbb{F}_2^{k\times n}$
  $\mathbf{L}^{(2,3)} \leftarrow \mathbf{I}_{n\times n}$ ▷ Identity matrix
  while zero_columns($\mathbf{B}$) $< n-k$ **do** ▷ Number of zero-columns is checked.
    $temp \leftarrow 0$
    for j=0; j<n; j++ **do**
      for i=0; i<n; i++ **do**
        if $j \neq i$ and $2w(\mathbf{b}_{:,j}) > w(\mathbf{b}_{:,i})$ **then**
          if $w(\mathbf{b}_{:,j} \oplus \mathbf{b}_{:,i}) < w(\mathbf{b}_{:,j})$ **then**
            $\mathbf{b}_{:,j} \leftarrow \mathbf{b}_{:,j} \oplus \mathbf{b}_{:,i}$
            $\mathbf{l}_{:,j}^{(2,3)} \leftarrow \mathbf{l}_{:,j}^{(2,3)} \oplus \mathbf{l}_{:,i}^{(2,3)}$
            $temp \leftarrow 1$
    if temp $= 0$ **then**
      break;
  return $\mathbf{B}, \mathbf{L}^{(2,3)}$ ▷ $\mathbf{B}$:Check type, $\mathbf{L}^{(2,3)}$:Local sets.

equations. For example, suppose we have the following check #1 local group defined for data nodes D_0, D_1 and D_2: (D_0, D_1, D_2). Also, suppose that we have the following check #2 equations:

· $D_0 = C_0 \oplus C_1 \oplus C_2$

· $D_1 = C_1 \oplus C_{12} \oplus C_{17} \oplus C_{20} \oplus C_{99}$

· $D_2 = C_0 \oplus C_{21} \oplus C_{99}$

Thus, we can find a check #3 type equation given by $(C_2, C_{12}, C_{17}, C_{20}, C_{21})$ by observing the following equivalence,

$$C_2 \oplus C_{12} \oplus C_{17} \oplus C_{20} \oplus C_{21} = D_0 \oplus D_1 \oplus D_2 \tag{1}$$

Note that since this technique uses check #1, it is likely to generate distinct check #3 local groups and help improve repair performance.

2.6. Algorithm for generating Group #2 and Group #3 check equations

We use a heuristic algorithm to generate Group #2 and Group #3 check equations at the same time. We let $\mathbf{B} \in \mathbb{F}_2^{k\times n}$ be the input generator matrix of the Founsure base graph code. We use $\mathbf{b}_{i,:}$ to refer to the ith row of $\mathbf{B}$ and $\mathbf{b}_{:,i}$ to refer to the ith column of $\mathbf{B}$. The Algorithm uses XOR operation ($\oplus$) to sparsify the generator matrix $\mathbf{B}$. If $\mathbf{B}$ is full rank, i.e., $rank(\mathbf{B}) = k$ and algorithm converges succesfully, at the end we should be able to find a permutation matrix $\mathbf{P}$ such that $\mathbf{BP} = [\mathbf{I}_{k\times k} \mid \mathbf{0}_{k\times n-k}]$. Also, $\tilde{\mathbf{L}}^{(2,3)} = \mathbf{L}^{(2,3)}\mathbf{P}$ shall hold all the local sets i.e., Check #2 equations in the first k columns and Check #3 equations in the last $n-k$ columns. We can express different types of checks as the union of all Check #2 and Check #3 check equations as given by

$$\bigcup_{i<k}(D_i, \{C_s : \tilde{\mathbf{l}}_{s,i} \neq 0\}) \quad \cup \quad \bigcup_{i\geq k}(\{C_s : \tilde{\mathbf{l}}_{s,i} \neq 0\}) \tag{2}$$

Here we provide the details of the algorithm below using a pseudocode formulation. We use a simple function zero_columns(.) that finds the number of nonzero columns of the matrix in the argument. Since $\mathbf{B}$ is typically sparse, we use sparse representation of matrices in our Founsure 1.0 implementation and XOR operation is replaced by set-union operation.

2.7. Managing Check Equations for Group #2 and Group #3 and generation of <filename>_ check.data file

Check #1 is determined by a binary LDPC array codes as explained before. This choice is user specific and its rate usually depends on the precode rate required by the code itself as well as the reliability imposed by the application. The graph connections are given by the constraints of array codes.

Unlike Check Equations #1, Check equations #2 and #3 are determined by the Founsure base code i.e., LT-like (fountain) code. Based on the generator matrix of the code, an iterative heuristic low-complexity algorithm is run to determine n equations. If the algorithm converges, then we should have k Equations for #2 and $n-k$ Equations for #3. The algorithm produces a correct set of local equations (sets) but does not guarantee those equations to be independent.

However, in generating those equations, we do not employ any matrix inversion (which is quite costly for large size matrices) to find check equations and hence we trade off the efficiency by performance. The function that generates check #2 and #3 local groups is the utility function **genChecks**.

The function **genChecks** assumes that a metadata file is already generated by a previous run of the encoder **founsureEnc**. **genChecks** generates check groups and modifies the *meta_ data* file (appends the size of check data in terms of *sizeof(int)* bytes at the end of the metadata file if "-m" flag parameter is True). The check information is stored in another binary file called *<filename>_ check.data*. This file stores an integer array with a specific format as a binary file. The reason for introducing a format is to use bulk read/write capabilities of *fread* and *fwrite* C library functions which will make kernel's I/O performance acceptable.

The format is pretty straightforward and can be improved. Note that the integer value of the first *sizeof(int)* bytes is either 0 or 1.

- If it is 0 (Group #3), then the next integer value (next *sizeof(int)* bytes) gives the degree number i.e., the number of integers to be read as part of one local group for coded symbols.
- If it is 1 (Group #2), then the next integer value (next *sizeof(int)* bytes) gives the data symbol index which is involved with a local set whose degree is given by the following integer (next *sizeof(int)* bytes). This degree also indicates the next "degree" number, i.e., the number of integers to be read as part of one local group for coded symbols.

The beauty of the proposed **Algorithm 3** is that if it converges, then all of the data symbols are covered exactly by one particular Group #2 local set. Let us give an example and suppose that we have the following integer array stored in *<filename>_ check.data*:

0 4 13 56 17 66 1 19 2 11 13 0 2 39 88 ...

If we decode this integer array, we will be able to say that the first local set belongs to Group #3 and this set has four elements. In other words, 13th, 56th,17th and 66th coded symbols form a local group. The next local set belongs to Group #2 and the associated data symbol index is 19. This data symbol along with 11th and 13th coded symbols (two coded symbols) form a local group. This way we can decode the whole integer array stored in <filename>_check.data. If the algorithm converges, there should be k leading 1's and $n-k$ leading 0's in the integer array. Note that the total number of integers contained in the array is given by

$$N = \sum_{c=0}^{n-1} L_c + k + 2n \tag{3}$$

where L_c is the total number of elements in local group #2 indexed by c (without the data symbol) and Group #3 indexed by c. We also note that even if the algorithm does not converge,

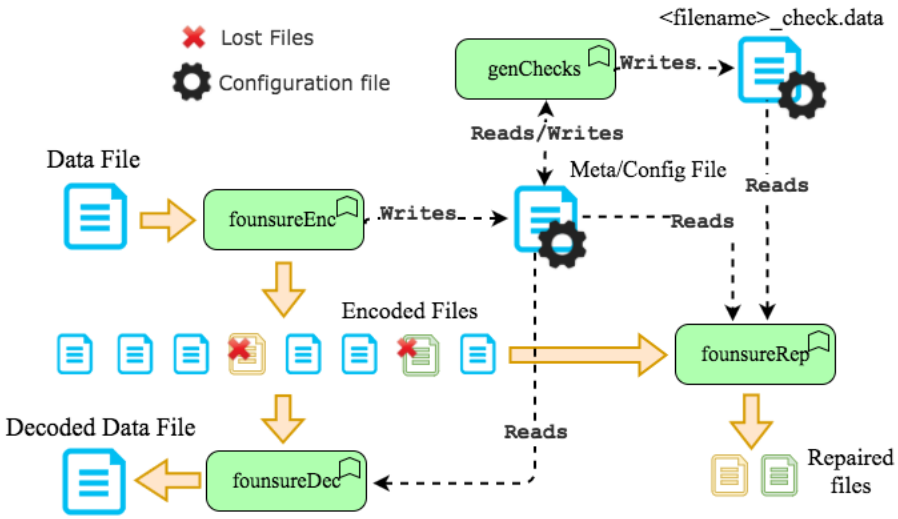

Figure 4: Summary of interactions between different functions of Founsure 1.0.

the maximum memory occupancy possible is $N\times$ *sizeof(int)* bytes, so it is enough to allocate that much memory for the file without encountering a segmentation fault.

Fig. 4 summarizes how different functions of Founsure interact with each other, what other metadata is generated and what are the read/write permissions granted to each of these functions.

2.8. Reading/Formating the contents of <filename>_ check.data file

When the repair process is initiated, memory allocation as well as repair object preparation starts. The main engine shall look for *<filename>_ check.data* under */Coding* directory. If it finds one and if the metadata is appropriately formatted, it will read-in and format check # 2 and check # 3 equations for the repair object. A bulk read is called and all the content is transferred to memory (inside the buffer pointed by *content2read*). Since Founsure's encoding, decoding, repair and update operations are based on Belief Propagation algorithm, it sequentially searches only one unknown over the available local sets in a loop. In order to reduce the computation and bandwidth, it is essential that repair/decode process use low-degree check # 3 equations first so that we do not have to run though the end of the loop to complete the process. Founsure implementation extracts check # 3 equations from the buffer (*content2read*) using the standard *qsort*(.) and then fills in the appropriate fields of repair object. The ordering can be enabled or disabled for check # 2 and # 3 equations using parameters ORDER_CHECK_2 and ORDER_CHECK_3 in "parameter.h" file.

2.9. Update Process

An update process is about making the code stronger or weaker by either generating more redundancy (in case of increased failures) or taking away unwanted redundancy (in case of using more reliable devices for storing information). If we would like to make the existent code weaker, it is easy. We just modify the metadata accordingly and erase the redundancy. So for the rest of this section, we particularly mean making the code stronger when we update the code.

Founsure update process is tightly related to repair process. This is mainly because updating a code is nothing but repairing the needed blocks of information to help increase the reliability of data. We call **genChecks** to update the current code using the flag '-e'. There must be a valid metadata associated with the code at the time **genChecks** is called. Code update process will rewrite n, the number of bytes used for the integer array due to check #2 and check #3 equations and update *<filename>_ check.data*. This series of modifications do not make any

changes to the data/coding chunks. In order to trigger/sync changes with data, we finally need to call **founsureRep** function with the appropriate file name.

3. Illustrative Examples

Here are the set of commands to use encoding, decoding, repair and update features of Founsure 1.0. Note that Founsure 1.0 comes with man pages or you can always use "-h" flag command for immediate help when you call Founsure functions.

The following command will encode a test file *testfile.txt* with $k = 500$ data chunks with each chunk occupying $t = 512$ bytes. The encoder generates $n = 1000$ coding chunks using $d =$'FiniteDist' degree distribution and $p =$'ArrayLDPC' precoding. Finally, generated chunks are striped/written to $s = 10$ distinct files for default disk/drive allocation under */Coding* directory[1]. The flag "-v" is used to output parameter information used during the encoding operation. Founsure encoder also generates a metadata file with critical coding parameters which will later be useful for decoding, repair and update operations. Without an appropriate metadata, Founsure cannot operate on files.

```
founsureEnc -f testfile.txt -k 500 -n 1000 -t 512 -d 'FiniteDist'
-p 'ArrayLDPC' -s 10 -v
```

Now, let us erase one of the coding chunks and run Founsure decoder. The decoder shall generate a decoded file *test_file_ decoded.txt* under */Coding* directory. You can use "diff" command to compare this file with the original.

```
rm -rf Coding/testfile_disk0007.txt
founsureDec -f testfile.txt -v
```

Note that **founsureDec** function does not recover the lost drive data *Coding/testfile_ disk0007.txt*, because this function is responsible only for the original data recovery process. In storage systems however, we need to recover lost data to keep the system data reliability at an acceptable level. In Founsure 1.0, it is extremely easy to initiate repair (current version only supports exact repair at the moment) process by running the following command.

```
founsureRep -f testfile.txt -v
```

This would trigger the conventional repair operation and first shall decode the entire data and then re-run partial encoding to generate the lost chunks. **founsureRep** outputs the pure computation speed as well as the bandwidth consumed due to repair. If you observe carefully, conventional repair is heavily time and bandwidth consuming operation. In fact, due to non-optimal overhead, the number of bytes that need to be transferred for the conventional repair is little larger than the size of the original user file. Founsure 1.0 supports fast and efficient repair as well. In order to use this feature, one needs to modify the metadata file and create an extra helping data/file called *testfile_ check.data* which shall contain information for fast repair. Details can be found later in the document. In order to make these changes, we primarily run **genChecks** function. Finally, we can re-run the repair function as before and you will realize from the comments printed out that the function will be able to recognize that there is available information for fast/efficient repair and will run that process instead of switching to conventional repair. You should be able to observe the reduced bandwidth consumed by the repair operation.

[1]This directory naming is conventional and maintained for the legacy of Jerasure erasure code library.

```
genChecks -f testfile.txt -m 1 -v
founsureRep -f testfile.txt -v
```

We can also use **genChecks** to trigger 'update' functionality. For example, let us assume that the system reliability is degraded due to drive wear and we want to generate extra two drive-worth information in addition to already generated 10 drive information. We use '-e' flag to modify metadata file as well as *testfile_ check.data* for the update operation. This shall change the code and all its related parameters. However in order to apply it to encoded data, we shall use **founsureRep** to generate new coding chunks and output files. Alternatively, you can erase drive info as well by supplying negative values for '-e' flag. In this case, you do not need to call **founsureRep** because there is nothing to generate. You can simply erase corresponding drive chunks after you scale the system down.

```
genChecks -f testfile -m 1 -v -e 2
founsureRep -f testfile.txt -v
```

4. Impact

To best of my knowledge, Founsure is the most flexible erasure coding library that is open source and can be configured based on the requirements of the application. With the current software architecture, many more functionalities can be integrated such as partial user data construction and advanced error detection for failure localization. Although the current version does not fully utilize, Founsure has highly parallel architecture and lends itself to parallel programming. Unlike Founsure, existing research is mostly focused on overhead optimal designs using the parallel hardware. Originally, Founsure is developped for data storage systems, it can simply be applied to packet switched networks in which the underlying channel is erasure channel or sporadic erasure channels. With advanced features such as error detection, multi-threaded implementations, advanced decoding, Founsure can further be used for error correction which would open up more fields of its application such as image reconstruction and data protection over noisy communication channels.

5. Conclusions

We developed an erasure coding library that can be used to solve the tradeoff between computation complexity, coding overhead and repair bandwidth. Founsure can be thought as the base software on which we can put lots of different features that will make it more application-centric and configurable for future generation reliable system requirements.

Acknowledgements

I would like to thank THE SCIENTIFIC AND TECHNOLOGICAL RESEARCH COUNCIL OF TURKEY (TUBITAK) and Quantum Corporation who provided the support for maturing ideas, as well as hardware platforms for extensive testing.

References

References

[1] B. Nisbet, "FAS Storage Systems: Laying the Foundation for Application Availability," Network Appliance white paper: http://www.netapp.com/us/library/analyst-reports/ar1056.html, February, 2008.

[2] I. S. Reed and G. Solomon, "Polynomial codes over certain finite fields," Journal of the Society for Industrial and Applied Mathematics, 8, 1960, pp. 300-304.

[3] J. Blomer, M. Kalfane, M. Karpinski, R. Karp, M. Luby and D. Zuckerman, "An XOR-Based Erasure-Resilient Coding Scheme," Technical Report TR-95-048, International Computer Science Institute, August, 1995.

[4] M. Blaum, J. Brady, J. Bruck and J. Menon, "EVENODD: An Efficient Scheme for Tolerating Double Disk Failures in RAID Architectures," IEEE Transactions on Computing, 44(2), February, 1995, pp. 192-202.

[5] M. Asteris, D. Papailiopoulos, A. G. Dimakis, R. Vadali, S. Chen, and D. Borthakur. Xoring elephants: Novel erasure codes for big data. Proceedings of the VLDB Endowment (PVLDB), 6(5), 2013.

[6] A. Partow, "Schifra Reed-Solomon ECC Library," Open source code distribution: http://www.schifra.com/downloads.html, 2000-2007.

[7] J. S. Plank, S. Simmerman and C. D. Schuman, "Jerasure: A Library in C/C++ Facilitating Erasure Coding for Storage Applications - Version 1.2," University of Tennessee, CS-08-627, August, 2008, http://www.cs.utk.edu/∼plank/plank/papers/CS-08-627.html.

[8] J. S. Plank, S. Simmerman, and C. D. Schuman, "Jerasure: A library in C/C++ facilitating erasure coding for storage applications - Version 1.2," University of Tennessee, Tech. Rep. CS-08-627, August 2008.

[9] Available online: http://www.ceph.com.

[10] S. S. Arslan, T. Goker and R. Wideman, "A Joint Dedupe-Fountain Coded Archival Storage," *IEEE ICC'17*, Paris, France, May 2017.

[11] J. Pearl. "Reverend Bayes on inference engines: A distributed hierarchical approach." *AAAI-82: Pittsburgh, PA. Second National Conference on Artificial Intelligence*, Menlo Park, California: AAAI Press. pp. 133–136. 1982.

[12] D. J. C. MacKay, "Fountain codes," in IEE Proc. Commun., vol. 152, pp. 1062–1068, Dec. 9, 2005.

[13] M. Luby, "LT-codes," in Proc. 43rd Annu. IEEE Symp. Foundations of Computer Science (FOCS), Vancouver, BC, Canada, Nov. 2002, pp. 271–280.

[14] A. Shokrollahi, "Raptor codes," IEEE Trans. on Information Theory, June 2006.

[15] Suayb S. Arslan "Incremental redundancy, fountain codes and advanced topics," 2014. Avaialable online: arXiv:1402.6016.

[16] A. J. Felstrom and K. S. Zigangirov, "Time-varying periodic convolutional codes with low-density parity-check matrix," IEEE Trans. Inf. Theory, vol. 45, no. 6, pp. 2181–2190, Sep. 1999.

[17] J. L. Fan, "Array codes as low-density parity-check codes," in Proc. 2nd Int'l Symposium on Turbo Codes and Related Topics, Brest, France, Sept. 2000, pp. 543-546,.

[18] E. Eleftheriou and S. Olcer, "Low-density parity-check codes for digital subscriber lines," in Proc. IEEE Int. Conf. Communications (ICC 2002), vol. 3, 2002, pp. 1752–1757.

Current code version

Nr.	Code metadata description	Please fill in this column
C1	Current code version	v10
C2	Permanent link to code/repository used for this code version	For example: *https* : *//github.com/suaybarslan/founsure*
C3	Legal Code License	LGPLv3
C4	Code versioning system used	git
C5	Software code languages, tools, and services used	C, Python, OpenMP, etc.
C6	Compilation requirements, operating environments & dependencies	ANSI C89 and C++98; Unix, Linux, MacOS, Windows
C7	If available Link to developer documentation/manual	For example: *http* : *//www.suaybarslan.com /founsure.html*
C8	Support email for questions	arslans@mef.edu.tr

Table 1: Code metadata (mandatory)